

Root games on Grassmannians

Kevin Purbhoo ^{*}
 University of British Columbia
 kevinp@math.ubc.ca

February 1, 2008

Abstract

We recall the *root game*, introduced in [P], which gives a fairly powerful sufficient condition for non-vanishing of Schubert calculus on a generalised flag manifold G/B . We show that it gives a *necessary and sufficient* rule for non-vanishing of Schubert calculus on Grassmannians. In particular, a Littlewood-Richardson number is non-zero if and only if it is possible to win the corresponding root game. More generally, the rule can be used to determine whether or not a product of several Schubert classes on $Gr_l(\mathbb{C}^n)$ is non-zero in a manifestly symmetric way. Finally, we give a geometric interpretation of root games for Grassmannian Schubert problems.

1 Prior work

In [P] we introduced the *root game*, a combinatorial game which can often determine whether or not a given Schubert structure constant is zero in the cohomology ring of a generalised flag manifold G/B . Our goal in this paper is to strengthen our earlier results in the case where the group G is $GL(n, \mathbb{C})$ and the Schubert classes are pulled back from a Grassmannian. We begin by recalling the root game for Schubert intersection numbers on the (ordinary) flag manifold $Fl(n)$. Apart from the root game, most of the relevant background material for this paper can be found in [F1, FG].

Let $G = GL(n)$. Let B and B_- denote the Borel subgroups of upper and lower triangular matrices respectively, and $T = B \cap B_-$ the standard maximal torus, consisting of invertible diagonal matrices.

Recall that for each element of the symmetric group $\pi \in S_n$, there is a corresponding T -fixed point πB on the flag manifold $Fl(n) = G/B$ (here we view π as an

^{*}Research partially supported by an NSERC scholarship.

element of $GL(n)$ via the standard representation of S_n), and an associated Schubert variety $X_\pi = \overline{B_- \cdot \pi B}$. We denote its cohomology class in $H^*(Fl(n))$ by $[X_\pi]$.

Our convention will be to write all permutations in one line notation

$$\pi = \pi(1)\pi(2)\dots\pi(n).$$

If $1 \in S_n$ denotes the identity element, and $w_0 = n \dots 321 \in S_n$ is the long word, then the Schubert class $[X_1]$ is the identity element in $H^*(G/B)$, and $[X_{w_0}] \in H^{\text{top}}(G/B)$ is the class of a point. In general $[X_\pi]$ is a class of degree $2\ell(\pi)$, where $\ell(\pi)$ denotes the length of π .

For $\pi_1, \dots, \pi_m \in S_n$, the Schubert intersection number

$$\int_{Fl(n)} [X_{\pi_1}] \cdots [X_{\pi_m}] \tag{1}$$

is always a non-negative integer. The root game attempts to determine whether this number is strictly positive.

The game is played on a set of squares $\mathcal{S} = \{S_{ij} \mid 1 \leq i < j \leq n\}$. In our diagrams, we will arrange the squares S_{ij} in an array, where i is the row index and j is the column index. In each square we allow tokens to appear. Each token has a label $k \in \{1, \dots, m\}$, and no square may ever contain two tokens with the same label. A token labelled k is called a k -token, and we write $k \in S_{ij}$ if a k -token appears in square S_{ij} .

A position in the game is specified by two pieces of data:

- The configuration of the tokens. Formally this is a map τ from \mathcal{S} to subsets of $\{1, \dots, m\}$, and our notation $k \in S$ is shorthand for $k \in \tau(S)$; however, in this paper we will wish to think of each token as a physical object which can be moved from square to square.
- A partition of the set of squares $\mathcal{S} = R_1 \sqcup \dots \sqcup R_r$. Each R_i is called a **region**.

The **initial position** of the game is as follows: there is a single region $R_1 = \mathcal{S}$, and for $i < j$, a k -token appears in square S_{ij} if and only if $\pi_k(i) > \pi_k(j)$.

From the initial position we move the tokens in the manner prescribed in the next paragraph. However, before each move we have the option of *splitting* regions into multiple regions. We define an **ideal subset** of the squares to be a set $A \subset \mathcal{S}$ with the property that if $S_{ij} \in A$, $i' \geq i$ and $j' \geq j$ then $S_{i'j'} \in A$. Given an ideal subset of the squares we can break up a region R into two regions: $R \cap A$ and $R \setminus A$. We call this **splitting R along A** , and we may repeat the process as many times as desired.

A **move** is specified by a region R , a token label $k \in \{1, \dots, m\}$ and a pair (i, j) with $1 \leq i < j \leq n$. After choosing these data, we move tokens as follows:

- For every h with $j < h \leq n$, if S_{jh} and S_{ih} are both in R and a k -token appears in S_{jh} but not in S_{ih} , we move the k -token from S_{jh} to S_{ih} ;

- For every h with $1 \leq h < i$, if S_{hi} and S_{hj} are both in R and a k -token appears in S_{hi} but not in S_{hj} , we move the k -token from S_{hi} to S_{hj} .

More succinctly put, within the region R we move k -tokens horizontally from column i to column j and vertically from row j to row i , wherever possible. See Figure 1.

In the play of the game we may make any sequence of moves in any order. The game is **won** when there is exactly one token in each square.

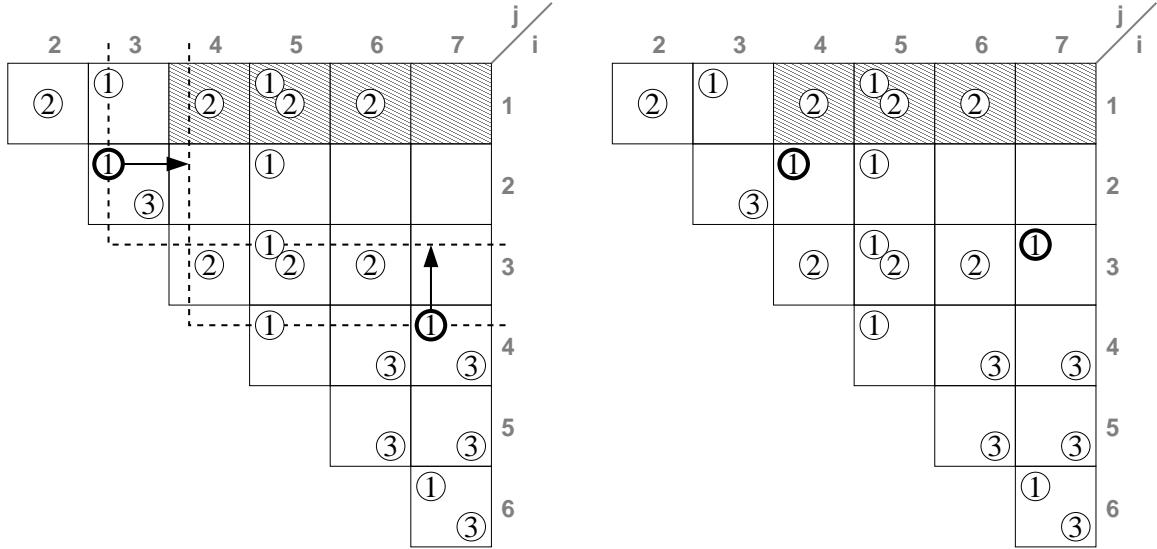


Figure 1: On the left is the initial position of the root game for $\pi_1 = 3426175$, $\pi_2 = 5162347$, $\pi_3 = 1326754$, after splitting into two regions. On the right is the position of the tokens after the move for $R =$ the unshaded region, $k = 1$, and $(i, j) = (3, 4)$. Note that the rows are indexed by $i = 1, \dots, 6$, and the columns by $j = 2, \dots, 7$, since $1 \leq i < j \leq 7$.

Remark 1.1. It turns out to be advantageous to split along an ideal subset A if and only if the total number of tokens in all squares in A equals $|A|$. We call the process of finding all such A and splitting along them **splitting maximally**. Although in this paper we won't take full advantage of this fact by splitting maximally before every move, we will never even consider the possibility of splitting along A if this condition is not met.

The main result that we shall need is the following.

Theorem 1 ([P]). *If the game for π_1, \dots, π_m can be won, then*

$$\int_{Fl(n)} [X_{\pi_1}] \cdots [X_{\pi_m}] \geq 1.$$

In general we do not know if the converse of Theorem 1 is true. When $m = 3$, the Schubert intersection numbers (1) are structure constants for the ring $H^*(Fl(n))$, and in this case the converse has been confirmed for $n \leq 7$. It would certainly be remarkable if it were true in general.

We can also use the game to study the cohomology rings of partial flag manifolds by pulling back cohomology classes to the full flag manifold. In this paper, we investigate this in some detail in the case of the Grassmannian. Our main result is a version of the converse of Theorem 1 for Grassmannian Schubert calculus.

When $m = 3$ and the classes come from a Grassmannian, the intersection numbers (1) are Littlewood-Richardson numbers. These numbers are also important in representation theory and in the theory of symmetric functions—they are the structure constants for the representation ring of $GL(n)$, and for the ring of symmetric functions in the Schur basis. As such, they are well studied, and there are a number of combinatorial rules, and geometric rules (e.g. [C, V]) known for computing these numbers.

Remark 1.2. An interesting and pleasant feature of root games is that the rules naturally extend to any number Schubert classes in a way which is manifestly symmetrical in these inputs: it is immediately clear from the definitions that reordering the input permutations π_1, \dots, π_m does not affect whether or not the game can be won. This property is not generally shared by other combinatorial or geometric rules for Littlewood-Richardson numbers. Most of the known rules have no manifest symmetry. Knutson-Tao puzzles [KTW] are perhaps the most manifestly symmetrical Littlewood-Richardson rule, having a 3-cyclic symmetry when $m = 3$, but this is lost when one attempts to generalise beyond triple intersections.

The manifest symmetry of the root game on $Fl(n)$, descends to the Grassmannian case. In Section 2 we partially break this symmetry, but as we explain in Section 4, the full symmetry is easily restored. To our knowledge, the only other manifestly symmetrical rule in Grassmannian Schubert calculus is the Horn recursion (see the survey article [F2]), which, like the root game, does not explicitly compute Littlewood-Richardson numbers and is only for determining which Schubert intersection numbers are strictly positive.

2 Associating a game to a Grassmannian Schubert calculus problem

Definition 2.1. A **01-string** is a string $\sigma = \sigma^1 \dots \sigma^n$ where each $\sigma^i \in \{0, 1\}$. A **$0^m 1^l$ -string** is a 01-string $\sigma = \sigma^1 \dots \sigma^{m+l}$, where exactly l of the σ^i are equal to 1.

Schubert varieties in the Grassmannian $Gr_l(n)$ are indexed by $0^{n-l} 1^l$ -strings. Fix a base flag

$$\{0\} = V_0 \subsetneq V_1 \subsetneq \dots \subsetneq V_n = \mathbb{C}^n$$

in \mathbb{C}^n . The Schubert variety $Y_\sigma \subset Gr_l(n)$ corresponding to σ is

$$Y_\sigma = \{y \subset \mathbb{C}^n \mid \dim y \cap V_i \geq \sigma^1 + \cdots + \sigma^i\}.$$

We denote its cohomology class in $H^*(Gr_l(n))$ by $[Y_\sigma]$. According to these conventions $[Y_{0\dots 01\dots 1}]$ is the identity element in $H^*(Gr_l(n))$ and $[Y_{1\dots 10\dots 0}] \in H^{\text{top}}(Gr_l(n))$ is the class of a point.

Given a list of $s+2$ $0^{n-l}1^l$ -strings $\sigma_1, \dots, \sigma_s, \mu, \nu$, we will wish to study the Grassmannian Schubert intersection numbers

$$\int_{Gr_l(n)} [Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu].$$

We do so by investigating an equivalent problem on a full flag manifold. The most obvious way to do this is to simply consider the product of the classes $\alpha^*([Y_{\sigma_i}])$ etc., under the natural map $\alpha : Fl(n) \rightarrow Gr_l(n)$, as in Lemma 2.2 below. However, our purposes require that we do things in a somewhat less straightforward way.

Let $N \geq 0$ be an integer, and let σ be a $0^{n-l}1^l$ -string. Let $i_1 < \cdots < i_{n-l}$ denote the positions of the zeroes in σ , and $j_1 < \cdots < j_l$ denote the positions of the ones. We define three ways to associate a permutation to the 01-string σ :

$$\begin{aligned} \pi(\sigma, N) &= i_1 \dots i_{n-l} j_1 \dots j_l (n+1)(n+2) \dots (n+N) \\ \pi'(\sigma, N) &= (i_1+N) \dots (i_{n-l}+N) 12 \dots N (j_1+N) \dots (j_l+N) \\ \pi''(\sigma, N) &= i_{n-l} \dots i_1 (n+N) \dots (n+1) j_l \dots j_1. \end{aligned}$$

From $\sigma_1, \dots, \sigma_s, \mu, \nu$ we produce a list of permutations, $\pi_1, \dots, \pi_{s+2} \in S_{n+N}$:

$$\begin{aligned} \pi_1 &= \pi(\sigma_1, N) \\ &\vdots \\ \pi_s &= \pi(\sigma_s, N) \\ \pi_{s+1} &= \pi'(\mu, N) \\ \pi_{s+2} &= \pi''(\nu, N). \end{aligned}$$

Proposition 2.1.

$$\int_{Gr_l(n)} [Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] = \int_{Fl(n+N)} [X_{\pi_1}] \cdots [X_{\pi_{s+2}}].$$

The proof is based on the following standard pullback calculations, whose proofs we omit.

Lemma 2.2. *Let $\alpha : Fl(n) \rightarrow Gr_l(n)$ be the map which forgets all but the l -dimensional subspace of the flag. Then $\alpha^*([Y_\sigma]) = [X_{\pi(\sigma,0)}]$.*

For a 01-string σ , let σ_+ denote the string σ followed by N ones, and let $+\sigma$ denote the string σ preceded by N ones.

Lemma 2.3. *Let $\beta : Gr_l(n) \rightarrow Gr_{l+N}(n+N)$ be the map $V \mapsto V \times \mathbb{C}^N \subset \mathbb{C}^n \times \mathbb{C}^N$. Then $\beta^*([Y_{\sigma_+}]) = [Y_\sigma]$. If σ' is not of the form σ_+ for some $0^{n-l}1^l$ -string σ then $\beta^*([Y_{\sigma'}]) = 0$.*

If \mathcal{F} is a partial flag variety and $h \in H^*(\mathcal{F})$ is a Schubert class, let h^\vee denote the opposite Schubert class, i.e. the unique Schubert class such that $\int_{\mathcal{F}} h \cdot h^\vee = 1$. For example $[X_\pi]^\vee = [X_{w_0\pi}]$, and $[Y_\sigma]^\vee = [Y_{\sigma^{\text{rev}}}]$ where $\sigma^{\text{rev}} = \sigma^n \dots \sigma^1$ is σ reversed.

If $h_1, \dots, h_r \in H^*(\mathcal{F})$ are Schubert classes then the statement that

$$\int_{\mathcal{F}} h_1 \cdots h_r = c \tag{2}$$

is equivalent to the statement that

$$h_1 \cdots \widehat{h}_i \cdots h_r = c h_i^\vee + \cdots \tag{3}$$

in the Schubert basis. We'll call Equation (3) the h_i -special version of Equation (2).

Proof of Proposition 2.1. Consider the equation in $H^*(Gr_{l+N}(n+N))$:

$$\int_{Gr_{l+N}(n+N)} [Y_{\sigma_1+}] \cdots [Y_{\sigma_{s+}}] [Y_{+\mu}] [Y_{\nu+}] = c. \tag{4}$$

If we take the $[Y_{+\mu}]$ -special version of Equation (4) and pull it back to $Gr_l(n)$, we get (using Lemma 2.3) the $[Y_\mu]$ -special version of

$$\int_{Gr_l(n)} [Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] = c.$$

On the other hand, if we take the $[Y_{\nu+}]$ -special version of Equation (4) and pull it back to $Fl(n+N)$, we get (using Lemma 2.2) the $[X_{\pi_{s+2}}]$ -special version of

$$\int_{Fl(n+N)} [X_{\pi_1}] \cdots [X_{\pi_{s+2}}] = c.$$

□

3 Non-vanishing for Grassmannians

Let $\sigma_1, \dots, \sigma_s, \mu, \nu$ be $0^{n-l}1^l$ -strings (if $s = 1$, we'll write σ instead of σ_1). For any given $N \geq 0$, we associate permutations π_1, \dots, π_{s+2} as before. Our goal in this section is to prove the following theorem.

Theorem 2. *Take N suitably large ($N \geq l$ will always suffice). The root game corresponding to π_1, \dots, π_{s+2} can be won if and only if*

$$\int_{Gr_l(n)} [Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] \geq 1.$$

Moreover, only moves involving tokens labelled $1, \dots, s$ are required.

Remark 3.1. The sufficient condition $N \geq l$ is not a sharp bound. This fact raises a number of interesting questions, which will be discussed in Sections 4 and 5. In the meantime the reader should not be alarmed by examples which use smaller values of N .

We shall first consider what happens in the case where $s = 1$.

Recall the correspondence between 01-strings and Young diagrams. Our Young diagrams will be in the French convention (the rows are left justified and increase in length as we move down). If σ is a 01-string σ , let $r_i(\sigma)$ denote the number of ones before the i^{th} zero. We associate to σ the Young diagram $\lambda(\sigma)$ whose i^{th} row is $r_i(\sigma)$, where we are allowing the possibility that some rows may have length 0.

If λ is a Young diagram, let $N + \lambda$ denote the Young diagram obtained by adding N squares to each row of λ *including those rows which contain 0 squares*.

The initial positions of the 1-tokens are in the shape of the Young diagram $\lambda_1 = \lambda(\sigma)$. The initial positions of the 2-tokens are in the shape of a Young diagram $\lambda_2 = N + \lambda(\mu)$. The squares that do *not* contain a 3-token are also in the shape of a Young diagram, which we'll denote $\lambda_{\bar{3}}$; viewed upside down, $\lambda_{\bar{3}}$ is the complement to $\lambda(\nu)$ inside an $(n-l) \times (l+N)$ rectangle. The lower left corner of each of the Young diagrams $\lambda_1, \lambda_2, \lambda_{\bar{3}}$ is in the square $S_{n-l, n-l+1}$. See Figure 2 for an illustration of how these shapes are generated.

If $\lambda_2 \not\subseteq \lambda_{\bar{3}}$, then it is a basic fact that $[Y_\mu] [Y_\nu] = 0 \in H^*(Gr_l(n))$ (in fact this is a necessary and sufficient condition). Therefore, we may assume that no square contains both a 2-token and a 3-token. The squares which contain neither a 2-token nor a 3-token are empty squares—since N is suitably large the 1-tokens are all to the left of these squares—and are in the shape of a skew-diagram $\lambda_{\bar{3}}/\lambda_2$.

At the outset of the game, some immediate splitting can occur. We split in such a way that each square containing a 3-token becomes a 1-square region of its own. (For some choices of (σ, μ, ν) it may be possible to split beyond this, but our argument is slightly simplified if we elect not to.) The remaining squares are those of $\lambda_{\bar{3}}$, which form what we call the **big region**. The big region is the only region which is unsolved; naturally, therefore, this will be the region in which all moves take place.

It is worth taking a moment to note how tokens move within the big region. A priori, a move (i, j) will cause some k -tokens to move horizontally and others to move vertically. However, since the rows of the big region are indexed by $\{1, \dots, n-l\}$, and the columns are indexed by the disjoint set $\{n-l+1, \dots, n+N\}$, these cannot both happen. If $i < j \leq n-l$ then k -tokens will move vertically from row j to row i . If $n-l+1 \leq i < j$ then k -tokens move horizontally from column i to column j . No tokens move if $i \leq n-l < j$.

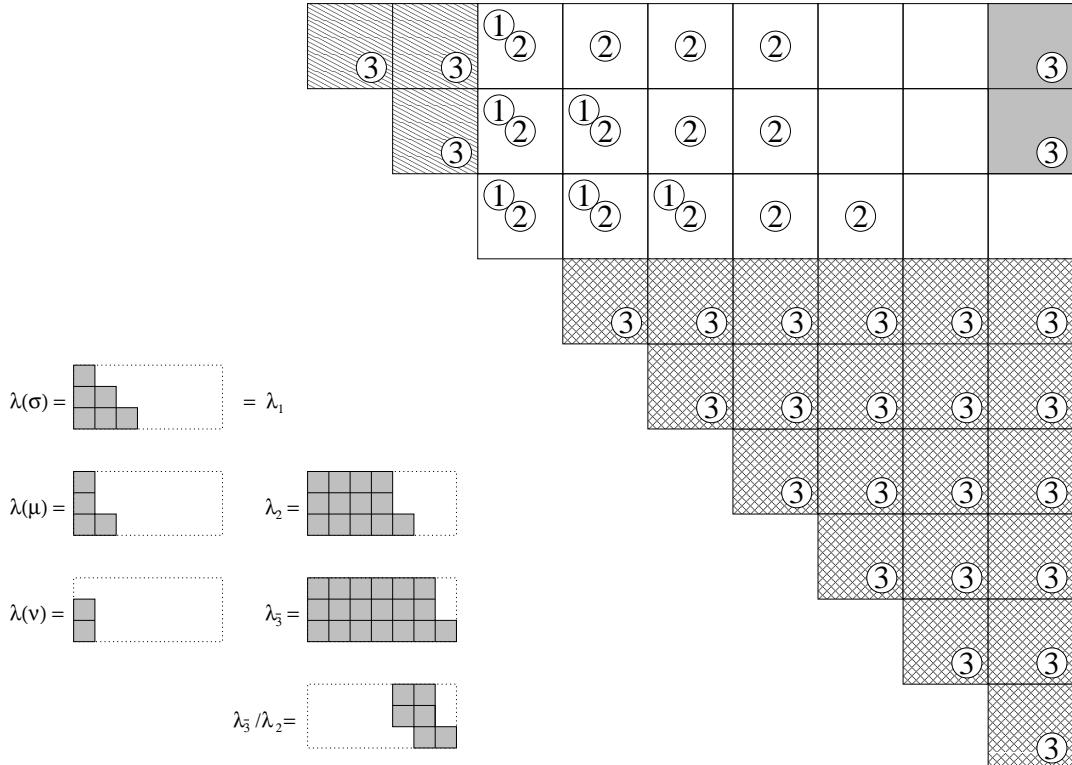


Figure 2: Initial position of the game for $\sigma = 1010101$, $\mu = 1001011$, $\nu = 0100111$, with $N = 3$. Here $\pi_1 = 24613578910$, $\pi_2 = 56812347910$, $\pi_3 = 43110987652$.

Definition 3.2 (Zelevinsky [Z]). A **picture** between two (French) skew diagrams is a bijection between their boxes with the property that if box A is weakly above and weakly right of box B in one diagram, then the corresponding boxes A' and B' are in lexicographic order (i.e. S_{ij} precedes $S_{i'j'}$ if $i < i'$ or $i = i'$ and $j < j'$) in the other diagram.

Now $\int_{Gr_l(n)} [Y_\sigma][Y_\mu][Y_\nu]$ is given by the Littlewood-Richardson coefficient $c_{\lambda_1 \lambda_2}^{\lambda_3}$, which can be described as the number of pictures between λ_1 and λ_3/λ_2 [Z]. This is a reformulation of the Littlewood-Richardson rule [LR], closely related to the Remmel-Whitney formula [RW]. In particular, if this number is non-zero, there exists such a

picture. We pick one, and denote by f the map it defines from the squares of λ_1 to the squares of $\lambda_{\bar{3}}/\lambda_2$. Note that all of these squares are in the big region of the game.

To show that the game can be won, we will give an algorithm—the *Grassmannian root game algorithm* (GRGA)—which uses f to construct a sequence of moves that transports each 1-token to a square of $\lambda_{\bar{3}}/\lambda_2$.

Essential to the GRGA is the following numbering scheme. At each point in the game we associate a number—called the *readiness number*—to each unplaced 1-token (i.e. one which has not already reached its final destination) and each empty square of $\lambda_{\bar{3}}/\lambda_2$.

Definition 3.3. At any point in the game, let t be a 1-token whose initial square was $S \in \lambda_1$ and whose current square is S_{ij} . Let $S_{i'j'} = f(S)$. Define the **readiness number** of both the token t and the square $S_{i'j'}$ to be the number $i - i'$. We say an empty square of $\lambda_{\bar{3}}/\lambda_2$ or an unplaced 1-token is **ready** if its readiness number is 0. Tokens which have reached their final destination and non-empty squares are *not* considered ready.

The key properties of this numbering scheme are the following:

Lemma 3.1. *Initially, the readiness numbers of the unplaced 1-tokens are*

- (a) *weakly increasing along each row, and*
- (b) *weakly decreasing down each column;*

in $\lambda_{\bar{3}}/\lambda_2$ the readiness numbers of empty squares are

- (c) *weakly decreasing along each row, and*
- (d) *weakly increasing down each column.*

Moreover, the GRGA, described below, preserves all of these properties.

Proof. (a) If A and B are squares of λ_1 in the same row, and A is right of B , then by Definition 3.2 $f(A)$ is lexicographically before $f(B)$. In particular, $f(A)$ is weakly above $f(B)$, i.e. in the same row or a row above. Thus the readiness number of the token in $B \leq$ the readiness number of the token in A .

(b) If A and B are squares of λ_1 , and A is one square above B , then again $f(A)$ is lexicographically before $f(B)$. There are two cases. If $f(A)$ is strictly above $f(B)$, then the readiness number of the token in $B \leq$ the readiness number of the token in A . Otherwise, $f(A)$ and $f(B)$ are in the same row in $\lambda_{\bar{3}}/\lambda_2$, with $f(B)$ right of $f(A)$. But then by Definition 3.2, B must be lexicographically before A , which is a contradiction.

Statements (c) and (d) are proved similarly. That the GRGA preserves all these properties will be quite evident. \square

Note that since the readiness number of the lower-leftmost token is at least 0, by Lemma 3.1 parts (a) and (b) the readiness numbers are initially all non-negative.

The Grassmannian root game algorithm (GRGA). *The algorithm assumes that the tokens are in the initial positions of the root game for π_1, π_2, π_3 (corresponding to σ, μ, ν with N suitably large), that all 3-tokens have been split into their own one-square region, and that we have a picture f between λ_1 and $\lambda_{\bar{3}}/\lambda_2$. All moves take place in the big region.*

1. If any of the 1-tokens are ready, go to Step 2. Otherwise, perform a sequence of moves to shift all unplaced 1-tokens up one square. The reader can easily check that the sequence of moves $(1, 2), (2, 3), \dots, (n-l-1, n-l)$ accomplishes this. The assumption that N is sufficiently large ensures that the upward movement of the 1-tokens is unobstructed. This step will decrease the readiness number of each 1-token by 1. Repeat this step until some 1-token is ready.
2. Scan through the columns of $\lambda_{\bar{3}}/\lambda_2$, beginning with the rightmost column and proceeding to the left. Within each column locate the topmost square that does not already contain a 1-token. Let S be the *first* ready square which we encounter in this way.
3. Find a ready token t in the same row as the square S . Make the unique move which causes t to move into S . This may cause other tokens to move as well. All tokens which move here move to their final destination, so after this move, they and the squares they occupy are no longer considered ready.
4. Repeat Steps 1 through 3 until every square of $\lambda_{\bar{3}}/\lambda_2$ contains a 1-token.

Example 3.4. Figures 3 and 4 illustrate the GRGA, applied to the example from Figure 2. We now draw only the squares in the upper right 3×7 rectangle as these are the only ones relevant to the movement of the 1-tokens. Moreover, only the 1-tokens are shown in these diagrams, and the number on the token is the readiness number, not the token label. To specify the picture f , each 1-token is given a shading and the corresponding square under f in $\lambda_{\bar{3}}/\lambda_2$ is shaded similarly. Each unshaded square actually contains a 2-token. The two darkly shaded squares in the upper right corner contain 3-tokens, as do each of the squares not shown in this diagram, but these squares are not part of the big region.

We now show that the GRGA accomplishes what it claims to accomplish.

Lemma 3.2. *Given a picture f between λ_1 and $\lambda_{\bar{3}}/\lambda_2$, the GRGA will win the game for π_1, π_2, π_3 .*

Proof. First note that after vertical movement from Step 1 is finished, the number of ready tokens in any row equals the number of ready squares in that row: an empty square S is ready if and only if the 1-token which began in the square $f^{-1}(S)$ is in the same row as S . We show that this equality is preserved, by showing that the move in Step 3 only ever causes ready tokens to move into ready squares.

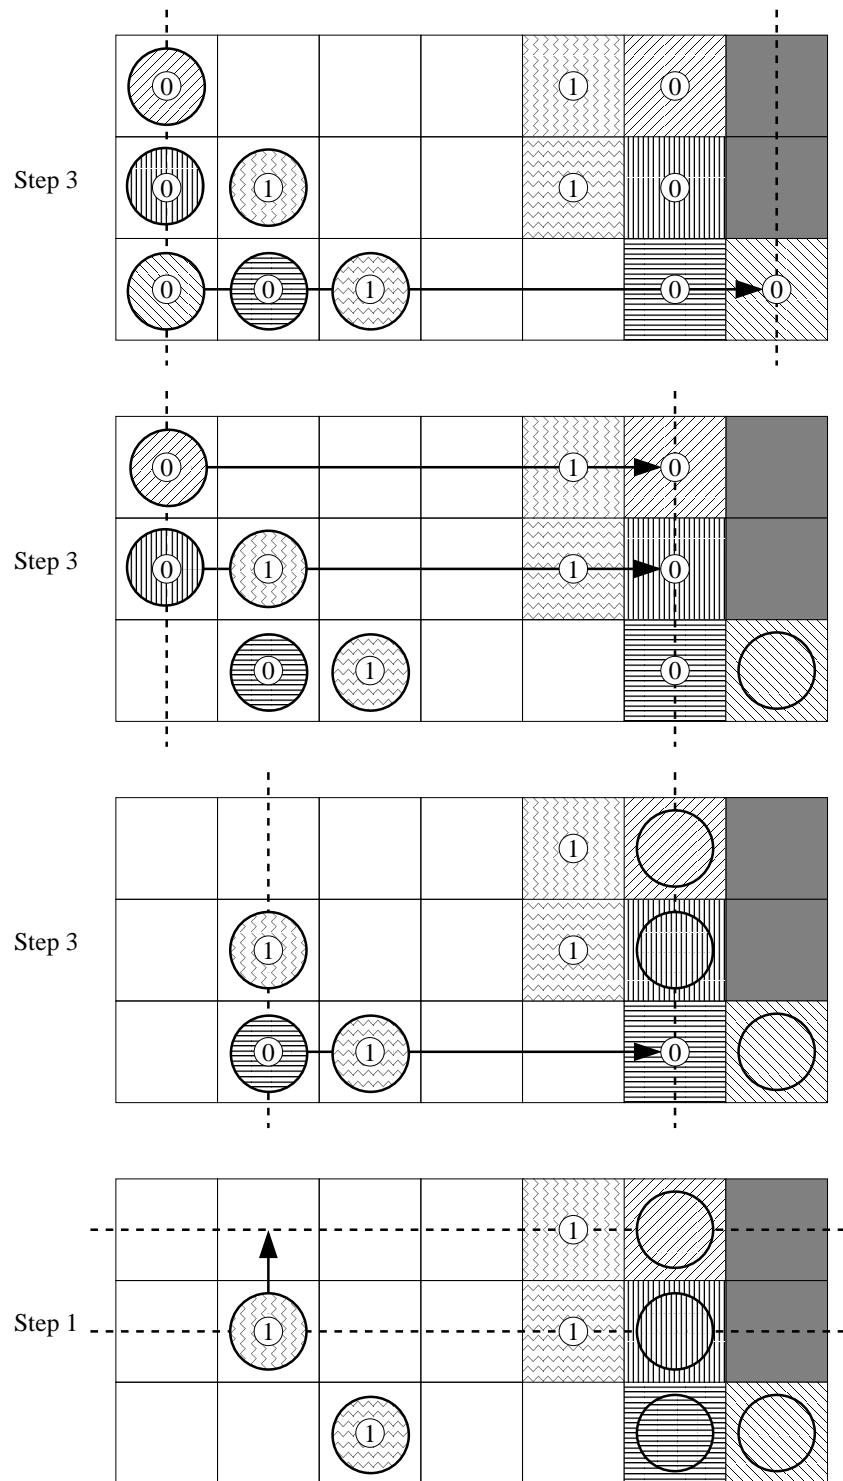


Figure 3: The Grassmannian root game algorithm. Here $\sigma = 1010101$, $\mu = 1001011$, $\nu = 0100111$, and $N = 3$.

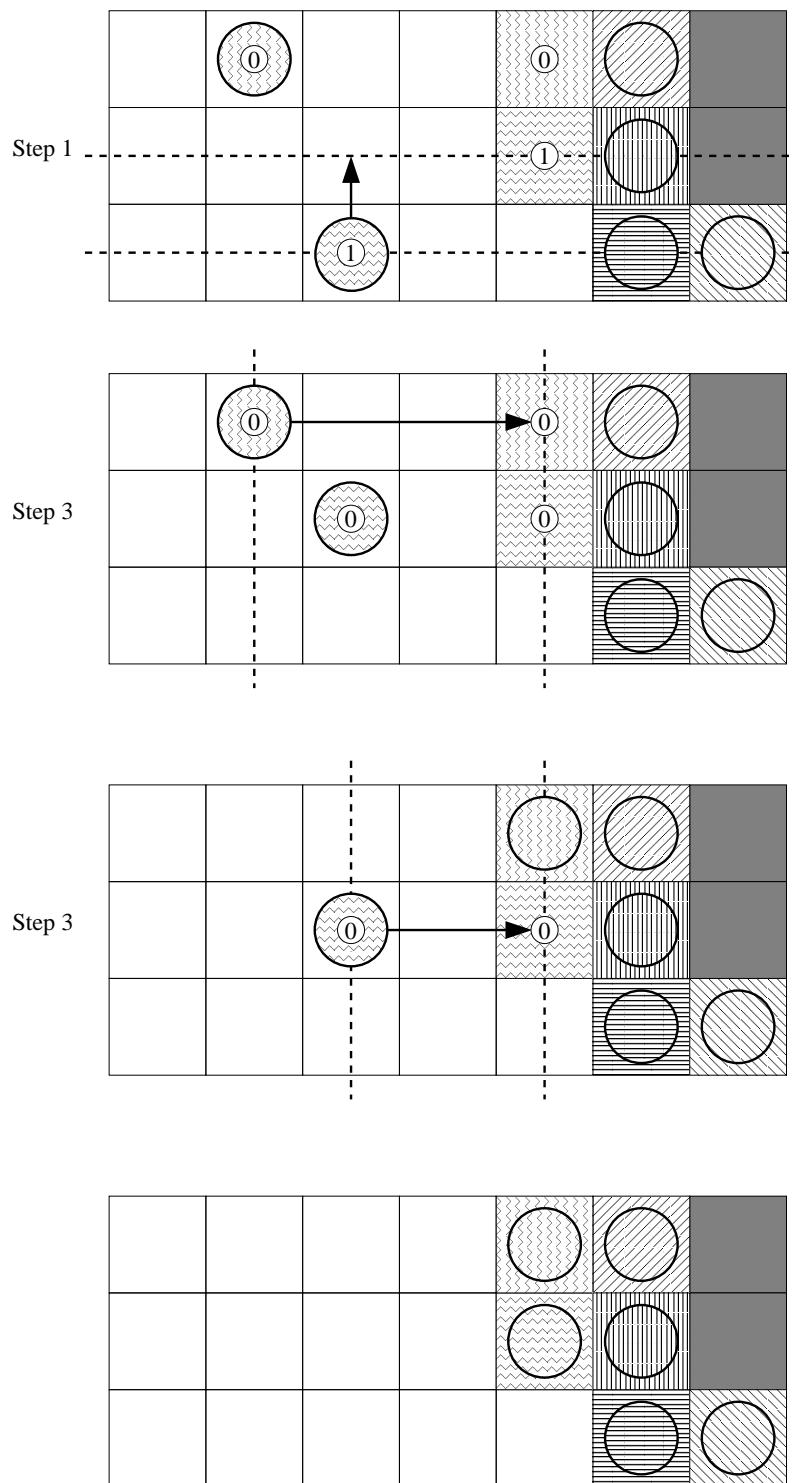


Figure 4: Continuation of Figure 3

Claim (i) *A move from Step 3 causes only ready tokens to move. In particular the number of ready tokens in a row always remains less than or equal to the number of ready squares in a row.* The only tokens that can conceivably move are those in the same column as t . Because S is the top unfilled square in its column, no tokens above t move. Because of Lemma 3.1 part (b), all tokens below t are ready.

Claim (ii) *Only ready squares are filled.* The algorithm attempts to fill the rightmost squares first. If there is a ready square S' in some column, the topmost empty square in that column will also be ready, by Lemma 3.1 part (d); thus the algorithm will never fill any square left of S' before it fills S' . However, by Lemma 3.1 part (c) the ready squares are rightmost in their row. Thus if a token moves into a non-ready square, it means that there are no ready squares in its row. But since only ready tokens move, we would have a row with at least one ready token and no ready squares. This, as noted in Claim (i), is impossible.

Claim (iii) *The move from Step 3 is always possible.* Since only ready squares are filled by ready tokens, the number of ready squares and ready tokens in any given row is always equal. Thus there is a ready token t in the same row as the ready square S . Because we assume N to be sufficiently large, t is to the left of S . If t is in column i and S is in column j , the move (i, j) will take the token t into square S .

Thus in Steps 2 and 3, every ready square eventually gets filled by a ready token: by the argument in Claim (ii) no square is skipped. However, because the readiness numbers are initially non-negative and Step 1 decreases the readiness number of each square by 1, every square of λ_3/λ_2 is ready at some point; thus the algorithm puts a 1-token in each square of λ_3/λ_2 , at which point the game is won. \square

Proof of Theorem 2. (\implies) This follows from Proposition 2.1 and Theorem 1.

(\impliedby) For $s = 1$ we use the GRGA, which wins the game by Lemma 3.2. For $s > 1$, we proceed by induction. Suppose $\int_{Fl(n+N)} [X_{\pi_1}] \cdots [X_{\pi_{s+2}}] \neq 0$. Then we can write

$$[X_{\pi_2}] \cdots [X_{\pi_{s+1}}] = c[X_\rho] + \cdots \quad (5)$$

in the Schubert basis, where $c > 0$, and

$$\int_{Fl(n+N)} [X_{\pi_1}] [X_\rho] [X_{\pi_{s+2}}] \neq 0.$$

Since this is really a Grassmannian calculation, ρ will be necessarily be of the form $\pi'(\sigma', N)$ for some $0^{n-l}1^l$ -string σ' . By Lemma 3.2 we can win the game corresponding to π_1, ρ, π_{s+2} , only moving 1-tokens. It is easy to see that exactly the same sequence of splittings and moves can be made in the game for $\pi_1, \pi_2, \dots, \pi_{s+1}, \pi_{s+2}$, and that it causes the 1-tokens to end up in exactly the same final positions. Note that we end up with either a 1-token or an $(s+2)$ -token in every square which does not correspond to an inversion of ρ , i.e. every square which does correspond to an inversion of $w_0\rho$.

This sequence of moves no longer wins the game; however, we can proceed inductively, after two further small steps. First, we perform a sequence of splittings so that every 1-token is in a one-square region of its own. Next, after splitting in this way, we replace each 1-token by an $(s+2)$ -token. This second step is not a legitimate play in the game, but it is completely harmless: because every 1-token is sequestered in its own one-square region, it can have no effect whatsoever on any possible subsequent moves of the game. But now we have precisely reached the initial position of the game corresponding to $\pi_2, \dots, \pi_{s+1}, w_0\rho$. This is again a game associated to a Grassmannian problem, and by Equation (5) the Schubert intersection number is non-zero. By induction, there is a sequence of moves to win this new game. Thus by concatenating the two sequences of moves, we can win the original game. \square

4 Remarks

In Step 3 of the GRGA, there is a somewhat canonical choice for the token t , namely the leftmost ready token in its row. If we use this choice of t , one can verify that the algorithm actually transports the 1-token which is initially in square S to the square $f(S)$. On the other hand, if there is a way of winning the root game, there is generally a plethora of ways, most of which do not arise by following the GRGA for any picture. Theorem 2 tells us that the existence of any one way to win implies the existence of a picture between λ_1 , and $\lambda_{\bar{3}}/\lambda_2$. However, given a sequence of moves which wins the game, it is not at all obvious how to construct such a picture. This is even unclear in the the case where $s = 1$, and only 1-tokens are moved.

It is worth noting that the root game can be used to determine whether

$$[Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] \neq 0$$

even if the cohomological degree of the product is not $\dim_{\mathbb{R}} Gr_l(n)$. To do this, we modify the game by changing the winning condition to read “the game is won if there is *at most* one token in each square”, rather than “exactly one token in each square”. Once we do this, we have the following corollary of Theorem 2.

Corollary 4.1. *Take N suitably large, and let π_1, \dots, π_{s+2} be obtained from $\sigma_1, \dots, \sigma_s$, μ, ν as before. Then $[Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] \neq 0$ if and only if the root game for π_1, \dots, π_{s+2} can be won with the modified winning condition.*

Proof. Assume $[Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] \neq 0$. Then there exists σ' such that

$$\int_{Gr_l(n)} [Y_{\sigma'}] [Y_{\sigma_1}] \cdots [Y_{\sigma_s}] [Y_\mu] [Y_\nu] \neq 0.$$

Let $\pi' = \pi(\sigma', N)$. Since we can win the unmodified game for $\pi', \pi_1, \dots, \pi_{s+2}$, we can win the modified game for π_1, \dots, π_{s+2} simply by omitting moves where the token corresponds to π' . The reverse direction follows from Proposition 2.1 and [P, Theorem 5] (which generalises Theorem 1). \square

There is a small catch: with this more general winning condition, our observation in Remark 1.1 becomes invalid. There is no longer an easy necessary and sufficient condition indicating when splitting is advantageous.

One of the unfortunate features of this presentation is the asymmetry in the way the permutations π_1, \dots, π_{s+2} are defined. The root game itself is manifestly symmetrical in the permutations given. However, because π_{s+1} and π_{s+2} are produced in a different way from π_1, \dots, π_s , the symmetry is broken for Grassmannians. Nevertheless, as Theorem 2 is valid for any s , we can formulate a symmetrical game by taking $\sigma_1, \dots, \sigma_s$ to be arbitrary, and $\nu = \mu = 0 \dots 01 \dots 1$, so that $[Y_\nu] = [Y_\mu] = 1 \in H^*(Gr_l(n))$. To see how this new formulation changes the initial position, contrast Figure 5 with Figure 2.

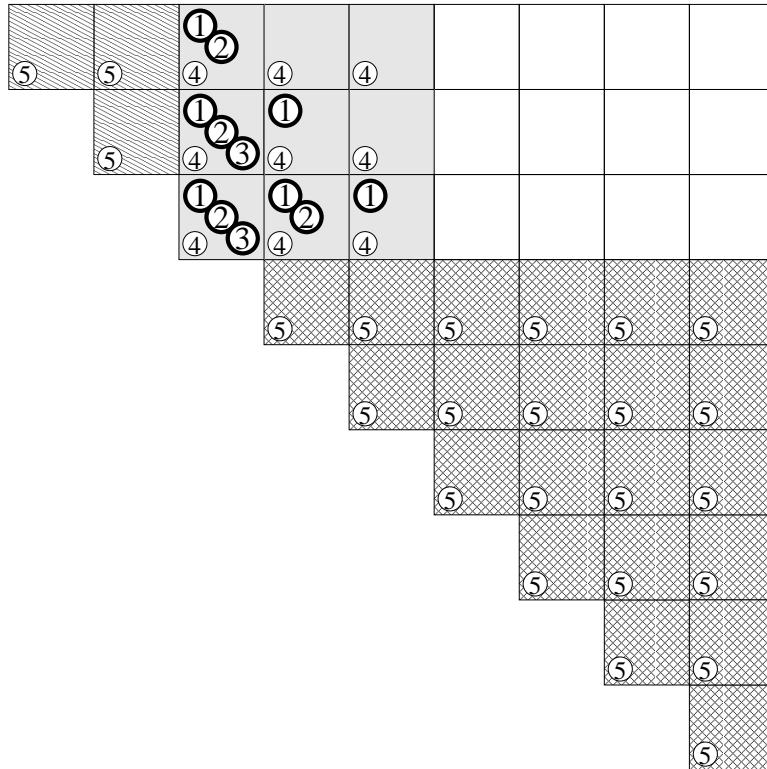


Figure 5: Initial position of the game for $\sigma_1 = 1010101$, $\sigma_2 = 1001011$, $\sigma_3 = 0100111$, $\mu = \nu = 0001111$, with $N = 3$. Squares are shaded if they contain a 4-token or a 5-token. This is the symmetrical version of the example in Figure 2.

The only splitting which occurs in the proof of Theorem 2 is before the first move, and in the inductive step. The GRGA itself does not split between moves. As noted in Remark 1.1, it can never be harmful to split maximally between moves, and it turns out that if one modifies the GRGA to split maximally between moves, things proceed very much as before. However, in the next section, our proof of Theorem 3

will rely on the fact that the GRGA involves no splitting.

It would be nice if we could take $N = 0$ in Theorem 2. Although we are not aware of any example which proves that this cannot be done, the algorithm simply falls apart if N is too small. There are several problems which occur with trying to follow a similar approach. The most serious of these is that a token may be to the right of the square for which it is supposedly destined according to the chosen picture. Again, this highlights the fact that we do not know a straightforward two-way correspondence between pictures and ways of winning the root-game. Instead, in the next section we prove Theorem 3, a geometrical analogue of Theorem 2 which is valid for all N . Theorem 3 suggests that it is not unreasonable to conjecture that Theorem 2 is true for all N . We leave it as an open problem to determine whether or not this is in fact the case.

5 Geometric interpretation

In [P] we give a complete description of the geometry underlying the root game. The picture is quite different from those found in the geometric Littlewood-Richardson rules of Vakil [V] and more recently Coskun [C], both of which study degenerations of intersections of Schubert varieties inside a Grassmannian—we would be surprised if there were any straightforward relationship. Our methods are based on studying tangent spaces to Schubert varieties, and are more closely related to the approach used by Belkale in his geometric proof of the Horn conjecture [B]. Here, we will recall only the parts of the picture which are most relevant to our situation.

In this section we shall once again assume $s = 1$. Our notation changes slightly from Section 1 in that we will be working with $GL(n+N)$ instead of $GL(n)$.

Let x_1, \dots, x_{n+N} denote the standard basis of \mathbb{C}^{n+N} . Let B denote the standard Borel subgroup of $GL(n+N)$ (upper triangular matrices), and let B_- denote its opposite (lower triangular matrices). As before $T = B \cap B_-$ will be the standard maximal torus. For any complex vector space V let $Gr(V)$ be the disjoint union of all Grassmannians $Gr_d(V)$, $0 \leq d \leq \dim V$.

Let R be a region in the game, and let

$$\tau : R \rightarrow \text{subsets of } \{1, 2, 3\}$$

describe the configuration of the tokens within this region. In the underlying geometry there is, assigned to the combinatorial pair (R, τ) , a corresponding geometric pair (V, U) , where V is a B -module, and $U = (U_1, U_2, U_3)$ is a T^3 -fixed point on $Gr(V)^3$ (or equivalently the U_k are T -invariant subspaces of V). As a T -representation, V is multiplicity-free, and the distinct T -weights are

$$\text{weights}(V) = \{x_i - x_j \mid S_{ij} \in R\}.$$

The (distinct) T -weights of U_k correspond to the positions of the k -tokens inside R :

$$\text{weights}(U_k) = \{x_i - x_j \mid k \in S_{ij}\}.$$

Thus the pair (V, U) carries all relevant information about the region R and the arrangement of the tokens with R .

The region R is solved when there is exactly one token in each square. In terms of the pair (V, U) this occurs when

$$V = U_1 \oplus U_2 \oplus U_3. \quad (6)$$

We'll call any $U = (U_1, U_2, U_3)$ which satisfies condition (6) **transverse**.

Assuming we do not split the region R , a move or a sequence of moves in the game takes the pair (V, U) to a new pair (V, U') , where U' is in the B^3 -orbit closure through $U \in \text{Gr}(V)^3$. Thus if we solve a region starting from position (V, U) , we have located a transverse T^3 -fixed point $U' \in \overline{B^3 \cdot U}$.

The importance of transverse points in $\overline{B^3 \cdot U}$ is seen in the following proposition.

Proposition 5.1 ([P]). *Consider the position of the root game game for π_1, π_2, π_3 which arises after splitting but before the first move is made. The tokens are in their initial position, but there may be more than one region. To each region R there is an associated pair (V_R, U_R) . Then*

$$\int_{Fl(n)} [X_{\pi_1}][X_{\pi_2}][X_{\pi_3}] \geq 1$$

if and only if for every region R there exists a transverse point $U' \in \overline{B^3 \cdot U_R}$.

Note that the point U' in Proposition 5.1 is not necessarily T^3 -fixed. The big question, therefore, is how specialised can we make the point U' and still have Proposition 5.1 be true. There are three levels of specialisation that we could request of this transverse point U' .

1. U' is any transverse point in $\overline{B^3 \cdot U_R}$.
2. U' is a T^3 -fixed transverse point in $\overline{B^3 \cdot U_R}$.
3. U' is a (T^3 -fixed) transverse point in $\overline{B^3 \cdot U_R}$, where (V_R, U') comes from applying sequence of root game moves (but no splitting) to the position (V_R, U_R) .

A priori, it is not clear that these three levels of specialisation are equivalent. However, for Grassmannian Schubert calculus with N sufficiently large, Theorem 2 shows that they are all equivalent. The GRGA tells us exactly how to produce a sequence of moves which gives the point U' at Level 3, which is the most specialised. Note it is important here that splitting is never used in the GRGA—when a region is split, the U' one is tempted to define need not be in $\overline{B^3 \cdot U_R}$.

Unfortunately, when N is small the GRGA can fail, and so we cannot claim that all three levels of specialisation are equivalent for all N . Our goal in this section is to show that even if N is too small for the GRGA to work, we can still get U' at

Level 2; i.e. Proposition 5.1 is still true for Grassmannian Schubert calculus if we demand that U' be a T^3 -fixed point.

To make matters more concrete, we now explicitly describe the initial pair (V, U) for the big region R in the root game associated to σ, μ, ν . This is the only region that we need to concern ourselves with, since all other regions are already solved.

Let $M(n+N)$ be the space of $(n+N) \times (n+N)$ matrices, having standard basis $\{\tilde{e}_{ij}\}$, and let B act on $M(n+N)$ by conjugation. Let W denote the B -submodule of $M(n+N)$ generated by the entries in the upper right $(n-l) \times (l+N)$ rectangle. Let W' be the B -submodule of W generated by \tilde{e}_{ij} such that S_{ij} contains a 3-token. Then V is the quotient B -module

$$V = W/W'.$$

Note V has a basis $\{e_{ij} := \tilde{e}_{ij} + W' \mid S_{ij} \in R\}$. The point $U \in Gr(V)^3$ is described as follows:

$$U_k = \text{span}\{e_{ij} \mid k \in S_{ij}\}.$$

Note that $U_3 = \{0\}$, so we need not give it much further consideration.

Let $U_{\overline{23}}$ be the subspace of V whose T -weights correspond to λ_3/λ_2 :

$$U_{\overline{23}} = \text{span}\{e_{ij} \mid 2 \notin S_{ij}\}$$

and let $U' = (U_{\overline{23}}, U_2, U_3)$. Note that $V = U_{\overline{23}} \oplus U_2 \oplus U_3$, so the point $U' \in Gr(V)^3$ is transverse.

Example 5.1. For the initial position shown in Figure 2,

$$\begin{aligned} V &= \text{span}\{\tilde{e}_{ij} \mid 1 \leq i \leq 3, 4 \leq j \leq 10\} / \text{span}\{\tilde{e}_{110}, \tilde{e}_{210}\} \\ &= \text{span}\{e_{14}, e_{15}, \dots, e_{19}, e_{24}, e_{25}, \dots, e_{29}, e_{34}, e_{35}, \dots, e_{39}, e_{310}\}, \\ U_1 &= \text{span}\{e_{14}, e_{24}, e_{25}, e_{34}, e_{35}, e_{36}\}, \\ U_2 &= \text{span}\{e_{14}, e_{15}, e_{16}, e_{17}, e_{24}, e_{25}, e_{26}, e_{27}, e_{34}, e_{35}, e_{36}, e_{37}, e_{38}\}, \\ U_{\overline{23}} &= \text{span}\{e_{18}, e_{19}, e_{28}, e_{29}, e_{39}, e_{310}\}. \end{aligned}$$

Theorem 3. For every $N \geq 0$,

$$\int_{Gr_l(n)} [Y_\sigma][Y_\mu][Y_\nu] \geq 1$$

if and only if with (V, U) and $U_{\overline{23}}$ as above, $U_{\overline{23}} \in \overline{B \cdot U_1}$.

Proof. (\Leftarrow) If $U_{\overline{23}} \in \overline{B \cdot U_1}$ then U' is transverse point in $\overline{B^3 \cdot U}$. Hence this follows from Propositions 2.1 and 5.1.

(\Rightarrow) Assume $\int_{Gr_l(n)} [Y_\sigma][Y_\mu][Y_\nu] \geq 1$.

We know the result is true for N sufficiently large, since the GRGA tells us how to get from the position (V, U) to the position (V, U') . We use this fact to deduce the result for other values of N .

For any two choices of N , say N^a and N^b , we get different spaces V , U , etc. We distinguish these notationally by using V^a (resp. V^b) to denote the space V corresponding to $N = N^a$ (resp. N^b), and likewise for any quantity depending on N . Note that $d = \dim U_1$ is independent of N .

For any fixed N^a and N^b , let $\phi : V^a \rightarrow V^b$ be the linear map given by

$$\phi(e_{ij}) = \begin{cases} e_{ij'} & \text{where } j' = j + N^b - N^a, S_{ij'} \in R^b \\ 0 & \text{if } S_{ij'} \notin R^b. \end{cases}$$

Let A be the dense open subset of $Gr_d(V^a)$,

$$A = \{X \in Gr_d(V^a) \mid X \cap \ker \phi = \{0\}\}.$$

Then ϕ induces a map $\phi_* : A \rightarrow Gr_d(V^b)$

$$\phi_*(X) = \text{Image } \phi|_X.$$

Observe that $\phi_*(U_{23}^a) = U_{23}^b$. The idea is essentially to show that $A \cap B^a \cdot U_1^a$ is dense in $\overline{B^a \cdot U_1^a}$, and that $\phi_*(A \cap B^a \cdot U_1^a) \subset \overline{B^b \cdot U_1^b}$. This implies that if $U_{23}^a \in \overline{B^a \cdot U_1^a}$, then $U_{23}^b = \phi_*(U_{23}^a) \in \overline{B^b \cdot U_1^b}$. Hence if the result is true for $N = N^a$, then the result will be true for $N = N^b$.

We will only prove this in the case where $N^a = 0$ and N^b is arbitrary, and in the case where N^a is arbitrary and $N^b = 0$. This is enough to give the result for all N .

The case where $N^a = 0$ is the easier of the two. The map

$$\begin{bmatrix} A^{n-l \times n-l} & B^{n-l \times l} \\ 0 & C^{l \times l} \end{bmatrix} \in B^a \mapsto \begin{bmatrix} A^{n-l \times n-l} & 0 & B^{n-l \times l} \\ 0 & I^{N^b \times N^b} & 0 \\ 0 & 0 & C^{l \times l} \end{bmatrix} \in B^b$$

allows us to view B^a as a subgroup of B^b , and ϕ_* is a B^a -equivariant inclusion. Moreover $\phi_*(U_1^a) \in \overline{B^b \cdot U_1^b}$. Thus ϕ_* takes $B^a \cdot U_1^a$ into $\overline{B^b \cdot U_1^b}$.

For the case where $N^b = 0$, we consider the B -orbit not through $U_1 \in Gr_d(V)$, but through a lifted point $\tilde{U}_1 \in Gr_d(W)$. \tilde{U}_1 is defined in the same way as U_1 : $\tilde{U}_1 = \text{span}\{\tilde{e}_{ij} \mid S_{ij} \text{ contains a 1-token}\}$. Let $\tilde{\phi}_*$ be defined analogously to ϕ_* , taking a dense subset of $Gr_d(W^a)$ to $Gr_d(W^b)$. It suffices to show that $\tilde{\phi}_*$ takes a dense subset of $B^a \cdot \tilde{U}_1^a$ to a subset of $B^b \cdot \tilde{U}_1^b$.

Let $L \cong GL(n-l) \times GL(l+N)$ be the subgroup of $GL(n+N)$ of block diagonal matrices of type $(n-l, l+N)$. Now L also acts on W , and \tilde{U}_1 is fixed by $B_- \cap L$. Since $(B \cap L) \cdot (B_- \cap L)$ is dense in L , it follows that the orbit $B \cdot \tilde{U}_1 = (B \cap L) \cdot \tilde{U}_1$ is dense in $L \cdot \tilde{U}_1$.

Thus in fact it suffices to show that ϕ_* takes a dense subset of $L^a \cdot \tilde{U}_1^a$ to a subset of $L^b \cdot \tilde{U}_1^b$. But this is true, as

$$\phi_* \left(\begin{bmatrix} A^{n-l \times n-l} & 0 & 0 \\ 0 & B^{l \times N^a} & C^{l \times l} \\ 0 & D^{N^a \times N^a} & E^{N^a \times l} \end{bmatrix} \cdot \tilde{U}_1^a \right) = \begin{bmatrix} A^{n-l \times n-l} & 0 \\ 0 & C^{l \times l} \end{bmatrix} \cdot \tilde{U}_1^b$$

whenever both matrices are invertible. □

Although Theorem 3 is a geometric statement, our proof ultimately relies on the combinatorics of the Littlewood-Richardson rule. The key non-geometric fact we use is that $\int_{Gr_l(n)} [Y_\sigma][Y_\mu][Y_\nu] \neq 0$ if and only if there exists a picture between λ_1 and $\lambda_{\bar{3}}/\lambda_2$. It would be an interesting project to find a purely geometric proof of this theorem. The hope would be that a geometric proof of Theorem 3 might allow us to see how to generalise some of the results in this paper beyond the Grassmannian.

Acknowledgements

The author is grateful to Allen Knutson and Stephanie van Willigenburg for providing feedback and corrections on this paper.

References

- [B] P. Belkale, *Geometric proofs of Horn and saturation conjectures*, J. Algebraic Geometry **15** (2006), no. 1, 133–173.
- [C] I. Coskun, *A Littlewood-Richardson rule for two-step flag manifolds*, preprint: <http://www-math.mit.edu/~coskun/>.
- [F1] W. Fulton, *Young tableaux with applications to representation theory and geometry*, Cambridge U.P., New York, 1997.
- [F2] W. Fulton, *Eigenvalues, invariant factors, highest weights, and Schubert calculus*, Bull. Amer. Math. Soc. **37** (2000), 209–250.
- [FG] S. Fomin and C. Greene, *A Littlewood-Richardson miscellany*, European J. Combin. **14** (1993), no. 3, 191–212.
- [KTW] A. Knutson, T. Tao, and C. Woodward, *The honeycomb model of $GL_n(\mathbb{C})$ tensor products II. Puzzles determine facets of the Littlewood-Richardson cone*, J. Amer. Math. Soc. **17** (2004), no. 1, 19–48 (electronic).
- [LR] D.E. Littlewood and A.R. Richardson, *Group characters and algebra*, Philos. Trans. Roy. Soc. London. **233** (1934), 99–141.
- [P] K. Purbhoo, *Vanishing and non-vanishing criteria in Schubert calculus*, Int. Math. Res. Not. (2006), 24590, 1–38.
- [RW] J.B. Remmel and R. Whitney, *Multiplying Schur functions*, J. of Algorithms **5** (1984), 471–487.

- [V] R. Vakil, *A geometric Littlewood-Richardson rule*, to appear in Ann. Math, arXiv preprint: [math.AG/0302294](https://arxiv.org/abs/math/0302294).
- [Z] A.V. Zelevinsky, *A generalization of the Littlewood-Richardson rule and the Robinson-Schensted-Knuth correspondence*, J. Algebra **69** (1981), no. 1, 82–94.